Equilibria of Infinitely Repeated Games

Game Theory Course:
Jackson, Leyton-Brown & Shoham

Strategy Space

- What is a pure strategy in an infinitely-repeated game?

Strategy Space

- What is a pure strategy in an infinitely-repeated game?
  - a choice of action at every decision point
  - here, that means an action at every stage game
  - ...which is an infinite number of actions!
- Some famous strategies (repeated PD):
  - Tit-for-tat: Start out cooperating. If the opponent defected, defect in the next round. Then go back to cooperation.
  - Trigger: Start out cooperating. If the opponent ever defects, defect forever.

Nash Equilibria

- With an infinite number of pure strategies, what can we say about Nash equilibria?
  - we won't be able to construct an induced normal form and then appeal to Nash's theorem to say that an equilibrium exists
  - Nash's theorem only applies to finite games
- Furthermore, with an infinite number of strategies, there could be an infinite number of pure-strategy equilibria!
- We can characterize a set of payoffs that are achievable under equilibrium, without having to enumerate the equilibria.

Definitions

- Consider any n-player game $G = (N, A, u)$ and any payoff vector $r = (r_1, r_2, \ldots, r_n)$.
- Let $v_i = \min_{s_{-i} \in S_{-i}} \max_{s_i \in S_i} u_i(s_{-i}, s_i)$.
  - i's minmax value: the amount of utility i can get when $-i$ play a minmax strategy against him

Definition

A payoff profile r is enforceable if $r_i \geq v_i$.

Definition

A payoff profile r is feasible if there exist rational, non-negative values α_a such that for all i, we can express r_i as $\sum_{a \in A} \alpha_a u_i(a)$, with $\sum_{a \in A} \alpha_a = 1$.

- feasible: a convex, rational combination of the outcomes in G.

Folk Theorem

Theorem (Folk Theorem)

Consider any n-player game G and any payoff vector $(r_1, r_2, \ldots, r_n)$.

1. *If r is the payoff in any Nash equilibrium of the infinitely repeated G with average rewards, then for each player i, r_i is enforceable.*
2. *If r is both feasible and enforceable, then r is the payoff in some Nash equilibrium of the infinitely repeated G with average rewards.*

Folk Theorem (Part 1)

Payoff in Nash $\Rightarrow$ enforceable

Part 1: Suppose r is not enforceable, i.e. $r_i < v_i$ for some i.

Folk Theorem (Part 1)

Payoff in Nash $\Rightarrow$ enforceable

Part 1: Suppose r is not enforceable, i.e. $r_i < v_i$ for some i. Then consider a deviation of this player i to $b_i(s_{-i}(h))$ for any history h of the repeated game, where b_i is any best-response action in the stage game and $s_{-i}(h)$ is the strategy of other players given the current history h.

Folk Theorem (Part 1)

Payoff in Nash $\Rightarrow$ enforceable

Part 1: Suppose r is not enforceable, i.e. $r_i < v_i$ for some i. Then consider a deviation of this player i to $b_i(s_{-i}(h))$ for any history h of the repeated game, where b_i is any best-response action in the stage game and $s_{-i}(h)$ is the strategy of other players given the current history h. By definition of a minmax strategy, player i will receive a payoff of at least v_i in every stage game if he adopts this strategy, and so i's average reward is also at least v_i. Thus i cannot receive the payoff $r_i < v_i$ in any Nash equilibrium.

Folk Theorem (Part 2)

Feasible and enforceable $\Rightarrow$ Nash

Part 2: Since r is a feasible payoff profile and the α's are rational, we can write it as $r_i = \sum_{a \in A} \left(\frac{\beta_a}{\gamma}\right) u_i(a)$, where β_a and γ are non-negative integers and $\gamma = \sum_{a \in A} \beta_a$.

Folk Theorem (Part 2)

Feasible and enforceable $\Rightarrow$ Nash

Part 2: Since r is a feasible payoff profile and the α's are rational, we can write it as $r_i = \sum_{a \in A} \left(\frac{\beta_a}{\gamma}\right) u_i(a)$, where β_a and γ are non-negative integers and $\gamma = \sum_{a \in A} \beta_a$.
We're going to construct a strategy profile that will cycle through all outcomes $a \in A$ of G with cycles of length γ, each cycle repeating action a exactly β_a times. Let (a^t) be such a sequence of outcomes.

Folk Theorem (Part 2)

Feasible and enforceable $\Rightarrow$ Nash

Part 2: Since r is a feasible payoff profile and the α's are rational, we can write it as $r_i = \sum_{a \in A} \left(\frac{\beta_a}{\gamma}\right) u_i(a)$, where β_a and γ are non-negative integers and $\gamma = \sum_{a \in A} \beta_a$.

We're going to construct a strategy profile that will cycle through all outcomes $a \in A$ of G with cycles of length γ, each cycle repeating action a exactly β_a times. Let (a^t) be such a sequence of outcomes. Let's define a strategy s_i of player i to be a trigger version of playing (a^t): if nobody deviates, then s_i plays a_i^t in period t.

Folk Theorem (Part 2)

Feasible and enforceable $\Rightarrow$ Nash

Part 2: Since r is a feasible payoff profile and the α's are rational, we can write it as $r_i = \sum_{a \in A} \left(\frac{\beta_a}{\gamma}\right) u_i(a)$, where β_a and γ are non-negative integers and $\gamma = \sum_{a \in A} \beta_a$.
We're going to construct a strategy profile that will cycle through all outcomes $a \in A$ of G with cycles of length γ, each cycle repeating action a exactly β_a times. Let (a^t) be such a sequence of outcomes. Let's define a strategy s_i of player i to be a trigger version of playing (a^t): if nobody deviates, then s_i plays a_i^t in period t. However, if there was a period t' in which some player $j \neq i$ deviated, then s_i will play $(p_{-j})_i$, where (p_{-j}) is a solution to the minimization problem in the definition of v_j.

Folk Theorem (Part 2)

Feasible and enforceable $\Rightarrow$ Nash

First observe that if everybody plays according to s_i, then, by construction, player i receives average payoff of r_i (look at averages over periods of length γ).

Folk Theorem (Part 2)

Feasible and enforceable ⇒ Nash

First observe that if everybody plays according to s_i, then, by construction, player i receives average payoff of r_i (look at averages over periods of length γ). Second, this strategy profile is a Nash equilibrium. Suppose everybody plays according to s_i, and player j deviates at some point.

Folk Theorem (Part 2)

Feasible and enforceable $\Rightarrow$ Nash

First observe that if everybody plays according to s_i, then, by construction, player i receives average payoff of r_i (look at averages over periods of length γ). Second, this strategy profile is a Nash equilibrium. Suppose everybody plays according to s_i, and player j deviates at some point. Then, forever after, player j will receive his $\min\max$ payoff $v_j \leq r_j$, rendering the deviation unprofitable.